APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
09/598,713	06/21/2000	Jonathan P. Douglas	2207/8613	7081

23838 7590 10/20/2003		EXAMINER
KENYON & KENYON 1500 K STREET, N.W., SUITE 700 WASHINGTON, DC 20005		HUISMAN, DAVID J

ART UNIT	PAPER NUMBER
2183	

DATE MAILED: 10/20/2003   5

Please find below and/or attached an Office communication concerning this application or proceeding.

	Application No.	Applicant(s)
● **Office Action Summary**	09/598,713 ●:	DOUGLAS, JONATHAN P. ᛞᛚ
	Examiner	Art Unit	
	David J. Huisman	2183	

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

> A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) FROM
> THE MAILING DATE OF THIS COMMUNICATION.
> - Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed
>   after SIX (6) MONTHS from the mailing date of this communication.
> - If the period for reply specified above is less than thirty (30) days, a reply within the statutory minimum of thirty (30) days will be considered timely.
> - If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
> - Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133).
> - Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any
>   earned patent term adjustment. See 37 CFR 1.704(b).

Status

1) ☒ Responsive to communication(s) filed on *15 August 2003* .

2a) ☐ This action is **FINAL**.   2b) ☒ This action is non-final.

3) ☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

4) ☒ Claim(s) *1,3-7 and 9-20* is/are pending in the application.

   4a) Of the above claim(s) _____ is/are withdrawn from consideration.

5) ☐ Claim(s) _____ is/are allowed.

6) ☒ Claim(s) *1,3-7 and 9-20* is/are rejected.

7) ☐ Claim(s) _____ is/are objected to.

8) ☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

9) ☒ The specification is objected to by the Examiner.

10) ☐ The drawing(s) filed on _____ is/are: a)☐ accepted or b)☐ objected to by the Examiner.

   Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).

11) ☒ The proposed drawing correction filed on *15 August 2003* is: a)☒ approved b)☐ disapproved by the Examiner.

   If approved, corrected drawings are required in reply to this Office action.

12) ☐ The oath or declaration is objected to by the Examiner.

Priority under 35 U.S.C. §§ 119 and 120

13) ☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).

   a)☐ All b)☐ Some * c)☐ None of:

   1. ☐ Certified copies of the priority documents have been received.

   2. ☐ Certified copies of the priority documents have been received in Application No. _____ .

   3. ☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).
   * See the attached detailed Office action for a list of the certified copies not received.

14) ☐ Acknowledgment is made of a claim for domestic priority under 35 U.S.C. § 119(e) (to a provisional application).

   a) ☐ The translation of the foreign language provisional application has been received.

15) ☐ Acknowledgment is made of a claim for domestic priority under 35 U.S.C. §§ 120 and/or 121.

Attachment(s)

1) ☒ Notice of References Cited (PTO-892)	4) ☐ Interview Summary (PTO-413) Paper No(s). _____ .
2) ☐ Notice of Draftsperson's Patent Drawing Review (PTO-948)	5) ☐ Notice of Informal Patent Application (PTO-152)
3) ☒ Information Disclosure Statement(s) (PTO-1449) Paper No(s) *4* .	6) ☐ Other: .

DETAILED ACTION

1.   Claims 1, 3-7, and 9-20 have been examined.

Papers Submitted

2.   It is hereby acknowledged that the following papers have been received and placed of

record in the file: #3. Amendment "A" as received on 8/15/2003 and #4. IDS as received on

8/15/2003.

Specification

3.   The abstract of the disclosure is objected to because it does not accurately describe the

contents of the disclosure. No mention is made of how the system deals with call and return

instructions, yet, the disclosure deals mostly with these concepts. Correction is required. See

MPEP § 608.01(b).

Claim Objections

4.   Claim 18 is objected to because of the following informalities: The examiner

recommends inserting --of cascaded pipestages-- after "the second plurality" in line 15. Also,

there is lack of antecedent basis for "the one pipestage" in line 17, since claim 17 refers to both a

first and second pipeline stages. Appropriate correction is required.

5.   Claim 19 is objected to because of the following informalities: The examiner

recommends inserting --of cascaded pipestages-- after "the first plurality" in line 21. Also, there

is lack of antecedent basis for "the one pipestage" in line 23, since claim 17 refers to both a first

and second pipeline stages. Appropriate correction is required.

Maintained Claim Rejections - 35 USC § 102

6.   The following is a quotation of the appropriate paragraphs of 35 U.S.C. 102 that form the

basis for the rejections under this section made in this Office action:

> A person shall be entitled to a patent unless –
>
> (b) the invention was patented or described in a printed publication in this or a foreign country or in public use or on sale in this country, more than one year prior to the date of application for patent in the United States.

7.   Claims 4 and 10-12 are rejected under 35 U.S.C. 102(b) as being anticipated by Hennessy

& Patterson, Computer Architecture - A Quantitative Approach, 2nd Edition, 1996 (as applied in

the previous Office Action and herein referred to as Hennessy).

8.   Referring to claim 4, Hennessy has taught an instruction pipe control method comprising:

a) reading a new instruction from an instruction pipestage. See page 154, Figure 3.13 and note

that each instruction is read from the appropriate stage. For instance, in clock cycle 1 (CC1), the

LW instruction is read from the fetch stage, in CC2, the LW instruction is read from the decode

stage, and so on.

b) determining, with reference to other instructions read previously from the instruction

pipestage, whether valid data associated with the new instruction can be written to a next

instruction pipestage. See page 154, Figure 3.13 and note, for instance, that if the SUB

instruction were to be read from the execution stage in CC4, a hazard would exist between the

SUB instruction and the previous LW instruction (since the SUB instruction is dependent on data

produced by the LW instruction). If no hazard were to exist, then the SUB instruction could be

executed in CC4 and valid data associated with the instruction could be written to a next pipeline stage.

c) stalling processing of the new instruction until valid data associated with the new instruction can be written to the next instruction pipestage. Since a hazard exists between the LW instruction and the SUB instruction, the SUB's associated data cannot be written to the next pipestage. As a result, the system inserts a bubble, or stalls for a cycle. Again, see Figure 3.13 on page 154.

d) if the new instruction is a call instruction, the determining includes determining whether immediate processing of the call instruction would exceed a predetermined access rate associated with a shared resource. Note that Hennessy has taught the basic concept of a call instruction on page 277. It should be further noted that each instruction within the instruction set travels through the processing pipeline, as shown on page 142, Figure 3.6. In this particular figure, if Instruction 3 were a call instruction, the immediate processing of the call would cause a structural hazard with the load instruction since both are accessing a memory, which is shared by the fetch and memory stages of the pipeline. In this case, the immediate processing of the call would exceed a predetermined access rate associated with the memory, where the access rate is one instruction per cycle.

9.   Referring to claim 10, Hennessy has taught a method for interfacing an instruction pipe with an external resource characterized by a predetermined round-trip communication latency period, the method comprising:

a) reading a new instruction from an instruction pipe stage. See page 154, Figure 3.13 and note that each instruction is read from the appropriate stage. For instance, in clock cycle 1 (CC1), the

LW instruction is read from the fetch stage, in CC2, the LW instruction is read from the decode

stage, and so on.

b) determining, with reference to other instructions read previously from the instruction

pipestage, whether valid data associated with the new instruction is available to the instruction

pipe prior to expiration of the round-trip communication latency period. It should be realized

that this constitutes basic dependency checking. For instance, on page 154, in Figure 3.12, a

RAW hazard exists between the LW instruction and the SUB instruction (i.e., the SUB

instruction needs to read an operand that is produced by the LW instruction). Therefore, since

the SUB is dependent on the LW, it will be determined that the data associated with the SUB is

not available prior to the expiration of the round-trip communication latency (that is, the latency

required for the load to access memory). On the other hand, if a SUB instruction immediately

following the LW is not dependent on the LW (and causes no other hazards), then it would be

determined that the data for that instruction would be available prior to the load finishing its

access of memory.

c) if not, stalling processing of the new instruction until the round-trip communication latency

period expires. If the data were not available, in the case of the RAW hazard, then the SUB

instruction would have to be stalled until the LW instruction finishes accessing memory. See

Figure 3.13 on page 154.

10.   Referring to claim 11, Hennessy has taught a method as described in claim 10. Hennessy

has further taught determining whether the new instruction requires access to the external

resource in excess of an access allocation for the instruction pipe, and if so, stalling the new

instruction. It should be noted that each instruction within the instruction set travels through the

processing pipeline, as shown on page 142, Figure 3.6. In this particular figure, if Instruction 3 were the new instruction, the immediate processing of the new instruction would cause a structural hazard with the load instruction since both are accessing a memory, which is shared by the fetch and memory stages of the pipeline. In this case, the immediate processing of Instruction 3 would exceed an access allocation for the instruction pipe, where the access allocation for the pipeline is one instruction per cycle is allowed to access the memory.

11.   Referring to claim 12, Hennessy has taught a method as described in claim 10. Hennessy has further taught that the stalling stalls the instruction pipestage and all other instruction pipestages before it in the instruction pipe. See Figure 3.13 on page 154 and note that if the SUB instruction is stalled before its 3^{rd} stage (as shown), then the next two subsequent instructions are stalled before their 2^{nd} and 1^{st} stages respectively.

Maintained Claim Rejections - 35 USC § 103

12.   The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

> (a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negatived by the manner in which the invention was made.

13.   Claims 5, 7, and 9 are rejected under 35 U.S.C. 103(a) as being unpatentable over Hennessy, as applied above, in view of Hoyt et al., U.S. Patent No. 5,604,877 (as applied in the previous Office Action and herein referred to as Hoyt).

14.   Referring to claim 5, Hennessy has taught an instruction pipe control method as described in claim 4. Hennessy has not explicitly taught after the stalling terminates, storing a return

address associated with the call instruction both locally and in a shared resource. However, Hoyt has taught a system for predicting return addresses wherein the prediction is made using either a register within the pipeline (Fig.5, component 45), or using a return stack buffer (Fig.5, component 51), which is normally implemented as a LIFO in main memory (column 2, lines 20-36). Hoyt has disclosed (in the Background section) that performing such prediction improves the efficiency of the system in that the processor can continue fetching instructions down the predicted path, thereby keeping the pipeline full. This advantage is also recognized by Hennessy (as shown in Figure 3.26 on page 167). Therefore, in order to aid in the prediction of return addresses, it would have been obvious to one of ordinary skill in the art at the time of the invention to modify Hennessy such that it is capable of predicting such addresses, as taught by Hoyt. And, part of this prediction system includes storing a return address both locally (in register 45, Fig.5) and in a shared resource (stack 51, Fig.5), where the return stack buffer is shared by components 40 and 60 (decoder and branch target buffer circuit). By storing the prediction locally, the prediction could be accessed more quickly as opposed to retrieving the prediction from the return stack buffer in memory. However, the return addresses must also be stored in the stack since the register can only hold one prediction. If there were no stack, then the system would lose return predictions if two or more calls happen in a row before a return. Therefore, it would have been obvious to one of ordinary skill in the art at the time of the invention to store the address both locally and in a shared resource.

15.   Referring to claim 7, Hennessy has taught an interface method for an instruction pipe that shares access to an external resource, comprising:

a) reading a new instruction from an instruction pipestage. See page 142, Figure 3.6 and note

that each instruction is read from the appropriate stage. For instance, in clock cycle 1 (CC1), the

LW instruction is read from the fetch stage, in CC2, the LW instruction is read from the decode

stage, and so on.

b) if the new instruction requires access to the external resource, determining with reference to

other instructions read previously from the instruction pipestage, whether immediate processing

of the new instruction would cause the instruction pipe to exceed an access allocation for the

instruction pipe. Again, see page 142, Figure 3.6. In this particular figure, if Instruction 3 were

a call instruction, the immediate processing of the call would cause a structural hazard with the

load instruction since both are accessing a memory, which is shared by the fetch and memory

stages of the pipeline. In this case, the immediate processing of the call would exceed a

predetermined access rate associated with the memory, where the access rate is one instruction

per cycle.

c) if so, stalling the new instruction. See page 143, Figure 3.7. Note that the structural hazard is

dealt with by stalling.

d) Hennessy has taught a call instruction (page 277) but has not taught that if the new instruction

is a call instruction, after the stalling terminates, storing a return address associated with the call

instruction both locally and in a shared resource. However, Hoyt has taught a system for

predicting return addresses wherein the prediction is made using either a register within the

pipeline (Fig.5, component 45), or using a return stack buffer (Fig.5, component 51), which is

normally implemented as a LIFO in main memory (column 2, lines 20-36). Hoyt has disclosed

(in the Background section) that performing such prediction improves the efficiency of the

system in that the processor can continue fetching instructions down the predicted path, thereby

keeping the pipeline full. This advantage is also recognized by Hennessy (as shown in Figure

3.26 on page 167). Therefore, in order to aid in the prediction of return addresses, it would have

been obvious to one of ordinary skill in the art at the time of the invention to modify Hennessy

such that it is capable of predicting such addresses, as taught by Hoyt. And, part of this

prediction system includes storing a return address both locally (in register 45, Fig.5) and in a

shared resource (stack 51, Fig.5), where the return stack buffer is shared by components 40 and

60 (decoder and branch target buffer circuit). By storing the prediction locally, the prediction

could be accessed more quickly as opposed to retrieving the prediction from the return stack

buffer in memory. However, the return addresses must also be stored in the stack since the

register can only hold one prediction. If there were no stack, then the system would lose return

predictions if two or more calls happen in a row before a return. Therefore, it would have been

obvious to one of ordinary skill in the art at the time of the invention to store the address both

locally and in a shared resource.

16.   Referring to claim 9, Hennessy in view of Hoyt has taught a method as described in claim

7. Hennessy has further taught that the stalling stalls the instruction pipestage and all other

instruction pipestages before it in the instruction pipe. See Figure 3.13 on page 154 and note that

if the SUB instruction is stalled before its 3^{rd} stage (as shown), then the next two subsequent

instructions are stalled before their 2^{nd} and 1^{st} stages respectively.

Claim Rejections - 35 USC § 102

17.   The following is a quotation of the appropriate paragraphs of 35 U.S.C. 102 that form the

basis for the rejections under this section made in this Office action:

> A person shall be entitled to a patent unless –
>
> (e) the invention was described in (1) an application for patent, published under section 122(b), by another filed in the United States before the invention by the applicant for patent or (2) a patent granted on an application for patent by another filed in the United States before the invention by the applicant for patent, except that an international application filed under the treaty defined in section 351(a) shall have the effects for purposes of this subsection of an application filed in the United States only if the international application designated the United States and was published under Article 21(2) of such treaty in the English language.

18.   Claims 13-14 are rejected under 35 U.S.C. 102(e) as being anticipated by Sproch et al.,

U.S. Patent No. 6,247,134 (herein referred to as Sproch).

19.   Referring to claim 13, Sproch has taught in an instruction pipe, a clock throttling

mechanism provided between a pair of instruction pipestages, comprising:

a) a state machine coupled to an output of a first instruction pipestage. See Fig.5 and note that

state machine 210 is coupled to the first stage of a pipeline. It determines whether to stall a

pipeline or not (2-state state machine).

b) a clock control circuit having an input for a system clock signal and having an output for a

modified clock signal, the output coupled to the pair of instruction pipestages, the clock control

circuit controlled by the state machine. See Fig.5 and see components 351, 352, 353, 355, 362,

363, and 365. These components represent the innards of component 230, shown in Fig.3. This

circuitry takes in a clock signal and modifies the output clock signals based on the state machine

circuit 210. Also, see column 8, lines 50-61.

20.   Referring to claim 14, Sproch has taught a mechanism as taught in claim 13. Sproch has

further taught a read/write controller under control of the state machine and having an output for

controlling writes to the second instruction pipestage. The circuit 230 shown in Fig.3, which

comprises the circuitry shown in Fig.5, as described above, also acts as a read/write controller,

which controls the reading and writing of data from one stage to the next. See column 8, lines

40-49.

Claim Rejections - 35 USC § 103

21.   The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all

obviousness rejections set forth in this Office action:

> (a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in
> section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are
> such that the subject matter as a whole would have been obvious at the time the invention was made to a person
> having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negatived by the
> manner in which the invention was made.

22.   Claims 1 and 6 are rejected under 35 U.S.C. 103(a) as being unpatentable over Hennessy,

as applied above, in view of Hoyt et al., U.S. Patent No. 5,604,877 (as applied in the previous

Office Action and herein referred to as Hoyt).

23.   Referring to claim 1, Hennessy has taught an instruction pipe control method comprising:

a) reading a new instruction from an instruction pipestage. See page 154, Figure 3.13 and note

that each instruction is read from the appropriate stage. For instance, in clock cycle 1 (CC1), the

LW instruction is read from the fetch stage, in CC2, the LW instruction is read from the decode

stage, and so on.

b) determining, with reference to other instructions read previously from the instruction

pipestage, whether valid data associated with the new instruction can be written to a next

instruction pipestage. See page 154, Figure 3.13 and note, for instance, that if the SUB

instruction were to be read from the execution stage in CC4, a hazard would exist between the

SUB instruction and the previous LW instruction (since the SUB instruction is dependent on data

produced by the LW instruction). If no hazard were to exist, then the SUB instruction could be

executed in CC4 and valid data associated with the instruction could be written to a next pipeline stage.

c) stalling processing of the new instruction until valid data associated with the new instruction can be written to the next instruction pipestage. Since a hazard exists between the LW instruction and the SUB instruction, the SUB's associated data cannot be written to the next pipestage. As a result, the system inserts a bubble, or stalls for a cycle. Again, see Figure 3.13 on page 154.

d) Hennessy has not taught that if the new instruction is a return instruction, the determining includes determining whether a return address is available within the instruction pipe. However, Hoyt has taught that if return address prediction is not used in a processor, then a pipeline will have to be stalled while the return instruction accesses main memory in order to retrieve the return address. As discussed by Hennessy, this is basically a type of hazard where the pipeline must stall until the required data is retrieved. If stalling did not occur, then the system would be at risk of producing incorrect results because the correct data would not be available for use. As a result, it would have been obvious to one of ordinary skill in the art at the time of the invention to stall when the new instruction is a return instruction (if the address were not available).

24.   Referring to claim 6, Hennessy has taught an instruction pipe control method as described in claim 1. Hennessy has further taught that the stalling stalls the instruction pipestage and all other instruction pipestages before it in the instruction pipe. See Figure 3.13 on page 154 and note that if the SUB instruction is stalled before its 3rd stage (as shown), then the next two subsequent instructions are stalled before their 2nd and 1st stages respectively.

25.   Claim 3 is rejected under 35 U.S.C. 103(a) as being unpatentable over Hennessy in view

of Hoyt, as applied above, and further in view of Skadron et al., Improving Prediction for

Procedure Returns with Return-Address-Stack Repair Mechanisms, 1998 (as applied in the

previous Office Action and herein referred to as Skadron).

26.   Referring to claim 3, Hennessy in view of Hoyt has taught an instruction pipe control

method as described in claim 1. Hennessy in view of Hoyt has not explicitly taught that if the

new instruction is a return instruction, the determining includes determining whether sufficient

time has expired from an earlier return instruction for a return address to be received from an

external resource. However, Skadron has shown concern for two return instructions that occur in

close succession (i.e., sufficient time has not expired from an earlier return before a second

return is encountered). See page 261, column 2, first full paragraph. Skadron has further

disclosed that this is a problem because if the first return cannot pop the stack before the second

one looks in the stack, then the second one reads from the wrong location. A person of ordinary

skill in the art would have recognized that in order to refrain from retrieving wrong predictions,

each return instruction should wait for a previous return instruction to finish. This is equivalent

to a structural hazard in Hennessy, which is defined on page 139 (of Hennessy). To summarize,

a resource conflict must result in a stall (in this case, the resource conflict would be multiple

returns accessing a memory stack). As a result, it would have been obvious to one of ordinary

skill in the art at the time of the invention to determine whether sufficient time has expired from

an earlier return instruction for a return address to be received from an external resource, in order

to ensure that the correct locations in the external resource are accessed.

27.   Claims 15-16 are rejected under 35 U.S.C. 103(a) as being unpatentable over Sproch, as

applied above, in view of Hennessy, as applied above.

28.   Referring to claim 15, Sproch has taught a mechanism as described in claim 13. Sproch

has not taught the specifics of claim 15. However, Hennessy has taught:

a) a first register coupled to the first instruction pipestage. See Figure 3.4 on page 134 and

Figure 3.5 on page 136, and note the ID/EX.A register (portion of ID/EX that is sent to the lower

input of the top MUX which is fed into the ALU). This register represents a result obtained

during the decode stage.

b) a second register. See Figure 3.4 on page 134 and Figure 3.5 on page 136, and note the

ID/EX.B register (portion of ID/EX that is sent to the upper input of the bottom MUX which is

fed into the ALU).

c) a selector coupled to the first and second registers and having an output coupled to the second

instruction pipestage. Note the MUXs that are connected to the register outputs. The MUXs

then have outputs which are ultimately fed into the EX/MEM register, which holds results

obtained during the EX stage.

A person of ordinary skill in the art would have recognized that Hennessy has taught a basic

pipeline configuration. It allows data to be routed to the appropriate portions of the processor,

thereby allowing instruction execution to occur efficiently. As a result, in order to execute

instructions and route the associated data to processor components, it would have been obvious

to one of ordinary skill in the art at the time of the invention to modify Sproch to include the

components taught by Hennessy.

29.   Referring to claim 16, Sproch in view of Hennessy has taught a mechanism as described

in claim 15. Furthermore, since the state machine controls the pipeline through circuitry 230

shown in Fig.3, and the selector is part of the pipeline as taught by Hennessy, then the selector is

also controlled by the state machine.

30.   Claims 17-20 are rejected under 35 U.S.C. 103(a) as being unpatentable over Pickett,

U.S. Patent No. 5,968,169, as applied above, in view of Sproch, as applied above.

31.   Referring to claim 17, Pickett has taught execution logic for a processor comprising:

a) a first instruction pipe, comprising a first plurality of cascaded pipestages. See Fig.15 and

note the cascaded stages of the pipeline (a new stage per clock cycle). The first pipe would

include one of the multiple decode units and multiple functional units in order to operate on one

many instructions fetched per clock cycle. See column 18, lines 31-38. Note that up to 6

instructions can be passed to 6 pipelines (Fig.1, components 208, 210, 212).

b) a return stack buffer (RSB) provided in communication with at least one of the first

pipestages. See Fig. 2 and note that the decode units from each pipeline will communicate with

the RSB.

c) a second instruction pipe, comprising:

    c1) a second plurality of cascaded pipestages, at least one of the second pipestages

    provided in communication with the return stack buffer. Note that a second pipeline

    would include a second set among the decode units, functional units, etc. This pipeline

    would also have pipestages as shown in Fig.15. Also, recall from Fig.2, that each

    pipeline's decode unit will communicate with the RSB.

c2) Pickett has not explicitly taught clock throttling logic couple to the at least one

second pipestage. However, Sproch has taught such a concept. See Fig.5. This circuitry

takes in a clock signal and modifies the output clock signals based on the need for a stall.

Also, see column 8, lines 50-61. The abstract of Sproch shows that such a concept allows

for power saving within the pipeline. Therefore, in order to save power, it would have

been obvious to one of ordinary skill in the art at the time of the invention to modify

Pickett to include clock throttling logic as taught by Sproch. It should also be noted from

Fig.5 that this logic is coupled to at least one pipestage of a pipeline.

32.   Referring to claim 18, Pickett in view of Sproch has taught logic as described in claim 17.

Sproch has further taught that the clock throttling logic comprises:

a) a state machine coupled to an output of the one pipestage from the second plurality. See Fig.5

and note that state machine 210 is coupled to the first stage of a pipeline. It determines whether

to stall a pipeline or not (2-state state machine).

b) a clock control circuit having an input for a system clock signal and having an output for a

modified clock signal, the output coupled to the one pipestage, the clock control circuit

controlled by the state machine. See Fig.5 and see components 351, 352, 353, 355, 362, 363,

and 365. These components represent the innards of component 230, shown in Fig.3. This

circuitry takes in a clock signal and modifies the output clock signals based on the state machine

circuit 210. Also, see column 8, lines 50-61.

33.   Referring to claim 19, Pickett in view of Sproch has taught logic as described in claim 17.

Furthermore, claim 19 is the same as claim 18 except that the state machine is coupled to the first

plurality in claim 19 as opposed to the second plurality in claim 18. However, one of ordinary

skill in the art would have recognized that if each pipeline were to include the ideas taught by

Sproch, then each pipeline would be able to save power, resulting in more overall system power

being saved. As a result, it would have been obvious to one of ordinary skill in the art at the time

of the invention to modify Pickett to have a state machine coupled to the first plurality.

34.   Referring to claim 20, Pickett in view of Sproch has taught logic as described in claim 17.

Pickett has further taught that additional instruction pipestages from either the first or the second

instruction pipe are provided in communication with the return stack buffer, the additional

instruction pipestages also provided with additional clock throttling logic. Recall from above

that the decode units (which operate in decode stages of pipelines) communicate with the RSB

via selector 258 as shown in Fig.2. In addition, as shown in Fig.2 and Fig.4, when a return

address is selected from the stack it is applied to the instruction cache as a fetch address.

Therefore, the RSB is also in communication with the fetch stage of the pipelines since a return

address is a fetch address. In addition, it should be noted from Sproch that the clock throttling

logic, which is used to save power, is coupled to each of the pipeline stages.

Response to Arguments

35.   In the remarks, Applicant argues the novelty/rejection of claim 4 on pages 8-9 of the

remarks, in substance that call instructions in Hennessy do not reference a shared memory

(instead they reference a branch target buffer (BTB)), and as a result, Hennessy has not taught

any stalling operation that occurs when processing a call instruction.

These arguments are not found persuasive for the following reasons:

a) The examiner agrees that call instructions generally reference a BTB. However, the issue here

is the first processing stage of the call instruction – the fetch stage. As shown in Fig.3.6 on page

142 of Hennessy, all instructions must be fetched. This is considered the first stage of processing

in this 5-stage pipeline. Therefore, when the call is to be fetched, it will generate a structural

hazard with the prior Load instruction (note the caption). This is because the load instruction is

accessing a memory to retrieve data and the memory also needs to be accessed in order to fetch

the call instruction. This will result in a hazard in a one-port memory, thereby resulting in a stall.

36.   In the remarks, Applicant argues the novelty/rejection of claim 7 on page 9 of the

remarks, in substance that Hennessy and Hoyt do not disclose storage of a return address in a

shared resource and that Hoyt never characterizes the stack 51 as being shared by any

components.

These arguments are not found persuasive for the following reasons:

a) It is clearly seen from Fig.5 of Hoyt that stack 51 (which is within branch address calculator

50) communicates with both the instruction decoder 60 and BTB circuit 40 (this is clear because

of the two-sided arrow connecting the components 60 and 50 (along with Fig.2, stage 2) and the

wire connecting 40 and 51). Therefore, components 50 and 51 are shared components. It should

also be realized that at the very least, component 50 is a shared component and stack 51 is a part

of the component 50. Therefore, storing an address in 51 is still storing it in component 50.

37.   In the remarks, Applicant argues the novelty/rejection of claim 10 on pages 9-10 of the remarks, in substance that the example cited in the Office Action is an example of data dependencies, not any round-trip communication latency.

These arguments are not found persuasive for the following reasons:

a) As described in the previous Office Action, the round-trip communication latency is the latency associated with accessing a memory. A load instruction, for instance, will make an access request to memory and information will be retrieved. Since memory has an access time associated with it, this communication (access and retrieval) will take some amount of time. This amount of time is the round-trip communication latency. Revisiting the rejection of claim 10, on page 154, in Figure 3.12, a RAW hazard exists between the LW instruction and the SUB instruction (i.e., the SUB instruction needs to read an operand that is produced by the LW instruction). Therefore, since the SUB is dependent on the LW, the SUB will not be able to progress in the pipeline if the data required by the SUB is not available. The data will become available at the expiration of the round-trip communication latency (that is, the latency required for the load to access memory and retrieve the data needed by the SUB).
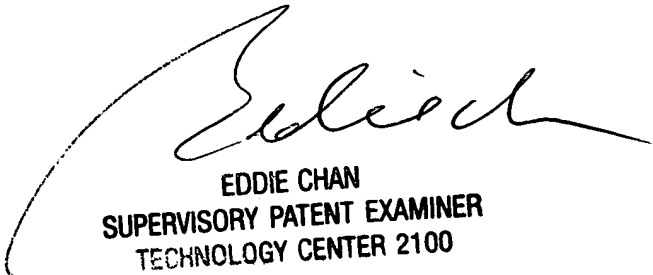
Conclusion

Any inquiry concerning this communication or earlier communications from the examiner should be directed to David J. Huisman whose telephone number is (703) 305-7811. The examiner can normally be reached on Monday-Friday (8:00-4:30).

If attempts to reach the examiner by telephone are unsuccessful, the examiner's

supervisor, Eddie Chan can be reached on (703) 305-9712. The fax phone number for the

organization where this application or proceeding is assigned is (703) 872-9306.

Any inquiry of a general nature or relating to the status of this application or proceeding

should be directed to the receptionist whose telephone number is (703) 305-3900.

DJH
David J. Huisman
October 14, 2003

EDDIE CHAN
SUPERVISORY PATENT EXAMINER
TECHNOLOGY CENTER 2100